

A gentle introduction to sheaves on graphs

Jakob Hansen

Abstract

This document provides an introduction to the language and theory of sheaves on graphs with an eye toward engineering and other applications. No familiarity with topology or commutative algebra is assumed. However, a basic familiarity with graphs, particularly spectral graph theory, and a working knowledge of linear algebra on an abstract level such as found in a second course in linear algebra are assumed.

1 What is a sheaf?

Sheaves are the canonical structure for attaching data to spaces. In algebraic topology and geometry, this can take many sophisticated, subtle forms, but we will deal with a simple version adapted to graphs and simplicial complexes. For our purposes, a graph is a collection of vertices and edges, together with an *incidence relation*, which tells us when a given vertex is incident to a given edge. In particular, we will allow edges to be incident to only one vertex. This may happen in two ways: A vertex may have a self-loop, or an edge may be open at one end without terminating in a vertex. A sheaf, then, is a way of assigning data to edges and vertices of a graph, with extra information telling us how the data over different parts of the graph is related.

Sheaves combine graphs and linear algebra, so we will also deal with vector spaces. For our purposes, vector spaces will be real vector spaces with an inner product, although nearly all of the theory will extend to complex vector spaces, and the parts of the theory that do not rely on an inner product extend to vector spaces over any field.

Definition 1.1 (Sheaves). Let G be a graph. A sheaf \mathcal{F} on G consists of a vector space $\mathcal{F}(v)$ for each vertex v of G , a vector space $\mathcal{F}(e)$ for each edge e of G , and a linear transformation $\mathcal{F}_{v \triangleleft e} : \mathcal{F}(v) \rightarrow \mathcal{F}(e)$ for each incident vertex-edge pair $v \triangleleft e$.




Figure 1: An illustration of the structure of a sheaf over a graph.

The vector spaces $\mathcal{F}(v)$ and $\mathcal{F}(e)$ are called the *stalks* of \mathcal{F} over v or e . The linear maps $\mathcal{F}_{v \triangleleft e}$ are sometimes called *restriction maps*. In agricultural terminology, a sheaf is a collection of stalks of grain bound together by twine; in mathematical terminology, a sheaf is a collection of stalks of data bound together by linear maps. The script letter \mathcal{F} is frequently used for sheaves; it is short for *faisceau*, the French word for sheaf.

The maps $\mathcal{F}_{v \triangleleft e}$ encode consistency requirements for our data. Whenever consistency can be verified locally, we have a sheaf structure. The simplest example of this is checking whether a real-valued function on the vertices of a connected graph is constant. This is a form of consistency that can be verified locally: one only needs to check that the function is constant across each edge. Sheaf theory vastly generalizes this sort of local consistency checking.

Definition 1.2 (Sections). Let \mathcal{F} be a sheaf on a graph G , and let W be a subset of the vertices of G . A *section* of \mathcal{F} over W is an choice of vectors $x_v \in \mathcal{F}(v)$ for each $v \in W$, such that whenever v and v' are incident to an edge e , we have $\mathcal{F}_{v \triangleleft e}(x_v) = \mathcal{F}_{v' \triangleleft e}(x_{v'})$.

Note that there is a natural way to add any two sections of \mathcal{F} over W , simply by adding their values in each stalk. Because the restriction maps are linear, the sum of two sections is again a section. Similarly, there is a natural scalar multiplication on sections. This means that the sections of \mathcal{F} over W form a vector space. This space may be denoted $\mathcal{F}(W)$, $\Gamma(W; \mathcal{F})$, or $H^0(W; \mathcal{F})$. The space $\Gamma(G; \mathcal{F})$ of sections of \mathcal{F} over all vertices in G is called the space of *global sections* of \mathcal{F} . Further, note that a global section

x_v uniquely determines values x_e on the edges, by $x_e = \mathcal{F}_{v \triangleleft e}(x_v) = \mathcal{F}_{v' \triangleleft e}(x_{v'})$.

Remark. The term *section* comes from “cross section.” One often thinks of a sheaf as having a horizontal portion given by the underlying graph, with a vertical dimension added from the stalks. A section consists of choosing a consistent copy of the graph within the stalks of the sheaf, or a cross section of the vertical part of the sheaf.

Example. Let V be a vector space. The *constant sheaf* \underline{V} on G is given by the data $\underline{V}(v) = V$ for all vertices v , $\underline{V}(e) = V$ for all edges e , and $\mathcal{F}_{v \triangleleft e} = I$ for all incident vertex-edge pairs (v, e) . For any subset of vertices W , $\Gamma(W; \underline{V})$ consists of the *locally constant* functions on W with values in V . (A function is locally constant if it is constant on each connected component of the subgraph induced by W .)

Example. As a subexample of the previous example, let $V = \mathbb{R}$. The constant sheaf $\underline{\mathbb{R}}$ plays a special role in the theory of sheaves. Its relationship to a sheaf \mathcal{F} determines the existence of global sections of \mathcal{F} .

Every sheaf has at least one global section: the constant section which is zero everywhere. (This is a consequence of the fact that the space of global sections of a sheaf is a vector space, and there is no empty vector space.) However, many sheaves have no other global sections.

Example. Consider the sheaf in Figure 2. Suppose there is a global section taking value $x(v_1)$ on v_1 . The condition to be a global section then means $x(v_2) = x(v_1)$ due to the upper left edge, and similarly $x(v_3) = x(v_1)$ due to the upper right edge. But the bottom edge forces $x(v_3) = -x(v_2)$, meaning $x(v_1) = -x(v_1)$, which can only be satisfied if all three values are zero.

In fact, having nontrivial global sections is a very special property. If G is a connected graph with at least one cycle and all stalks have the same dimension, a generic choice of restriction maps yields a sheaf with no global sections. One consequence of this fact is that if we take restriction maps from real-world measurements, our sheaf is vanishingly unlikely to have global sections, and we will need either to denoise our data, or to work with approximate global sections.

Definition 1.3 (Cochains). Let \mathcal{F} be a sheaf over a graph G . There are two spaces of *cochains* of \mathcal{F} . A zero-dimensional cochain (or 0-cochain) is a choice of a vector $x_v \in \mathcal{F}(v)$ for every vertex v ; a one-dimensional cochain (or 1-cochain) is a choice of $x_e \in \mathcal{F}(e)$ for every edge. The space of zero-dimensional cochains is denoted $C^0(G; \mathcal{F})$, and the space of one-dimensional cochains is denoted $C^1(G; \mathcal{F})$.




Figure 2: A sheaf with no nontrivial global sections.

In mathematical terms, the space of 0-cochains is the direct sum of the stalks over vertices, and the space of 1-cochains is the direct sum of the stalks over edges. That is,

$$C^0(G; \mathcal{F}) = \bigoplus_{v \in V} \mathcal{F}(v), \quad C^1(G; \mathcal{F}) = \bigoplus_{e \in E} \mathcal{F}(e).$$

The space of 0-cochains differs from the space of global sections because we do not impose any consistency requirements on our choices of vectors in each stalk. Therefore, the space of global sections is a subspace of the space of 0-cochains. In fact, the space of global sections is the kernel of a particular linear map from the space of 0-cochains to the space of 1-cochains. This map is called the *coboundary map*, and is defined as follows: First choose an orientation for the edges of G . This yields an *incidence number* $[v : e] \in \{\pm 1\}$ to each incidence vertex-edge pair $v \triangleleft e$ so that each edge has both a positive and a negative incidence number associated to it. (An edge with only one incident vertex can have either incidence number toward the vertex; the choice will make no difference in the sequel.) A choice of orientation for the graph G gives incidence numbers, by letting $[v : e] = 1$ if e is oriented away from v and $[v : e] = -1$ otherwise.

The coboundary map $\delta : C^0(G; \mathcal{F}) \rightarrow C^1(G; \mathcal{F})$ is given on vertex stalks by $\delta(x_v) = \sum_{v \triangleleft e} [v : e] \mathcal{F}_{v \triangleleft e}(x_v)$, and extends by linearity to all of $C^0(G; \mathcal{F})$. When necessary, we will disambiguate the sheaf to which δ pertains by a superscript: $\delta^{\mathcal{F}}$. We can also define δ by its output on each edge stalk: $(\delta x)_e = \mathcal{F}_{u \triangleleft e} x_u - \mathcal{F}_{v \triangleleft e} x_v$, where e is the edge from u to v .

Example. The coboundary map of the constant sheaf \mathbb{R} on G , given in matrix form, is the transpose of the (signed) incidence matrix of G .

Example. Consider again the sheaf in Figure 2. We can represent $C^0(G; \mathcal{F})$ by \mathbb{R}^3 , and we can also represent $C^1(G; \mathcal{F})$ by \mathbb{R}^3 . If we orient each edge so that it is negatively incident to the vertex with higher number, the coboundary map is given by the matrix

$$\delta = \begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \\ 0 & -1 & -1 \end{bmatrix}.$$

Note that this matrix has a trivial nullspace, which is due to the fact that this sheaf has no nontrivial global sections.

Remark. The term “cochain” comes from algebraic topology. One of its major tools, homology, deals with understanding “chains” in a topological space, or mappings of simplices into that space. For technical reasons, it is often useful to dualize this theory, producing cohomology, which studies cochains, or linear functionals on the space of chains. The cochains of sheaves on graphs are algebraically similar to these cohomological cochains, which is what provoked this use of terminology. Indeed, there is a way to define cohomology of sheaves, which was a major motivation for the development of the theory of sheaves.

This cohomological connection is what leads to the common terminology of the *zeroth cohomology* (or *degree-zero cohomology*) for $H^0(G; \mathcal{F}) = \ker \delta$. It may seem reasonable to consider not just the kernel of δ , but also its cokernel, i.e. the quotient space $C^1(G; \mathcal{F})/(\text{im } \delta)$, and indeed it is. This space is denoted $H^1(G; \mathcal{F})$, and is called the *first cohomology* of the sheaf \mathcal{F} . As a quotient space, its interpretation is typically less natural than that of the space of global sections of \mathcal{F} . However, the space $\text{im } \delta$ is analogous to the cut space of a graph. The cut space of a graph is the space of all real-valued functions on edges spanned by indicator functions of cut sets. A quick argument shows that this is equal to the space $\text{im } B^T$, where B is the incidence matrix of the graph. We can similarly interpret $\text{im } \delta$ as the space generated by “cuts” of the data communication structure described by \mathcal{F} over the graph.

The incidence matrix is most famous in spectral graph theory as a building block for the graph Laplacian: $L = BB^T$. In order to use spectral methods with sheaves, we build a Laplacian out of the coboundary map. This requires that we choose an inner product structure on the stalks of our sheaf, but in most applications, such an inner product is naturally forthcoming.

Definition 1.4 (Sheaf Laplacian). If \mathcal{F} is a sheaf on a graph G with coboundary map δ , the *sheaf Laplacian* of \mathcal{F} is $L_{\mathcal{F}} = \delta^* \delta$.

Proposition 1.1 (Structure of the sheaf Laplacian). *If \mathcal{F} is a sheaf on a graph G , the matrix of the sheaf Laplacian $L_{\mathcal{F}}$ has a symmetric block structure, with the columns divided into one partition for each vertex of G . If v_1 and v_2 are distinct and both incident to the edge e , then the entry in the (v_2, v_1) block is $-\mathcal{F}_{v_2 \triangleleft e}^* \mathcal{F}_{v_1 \triangleleft e}$, and the entry in the (v_1, v_1) block is $\sum_{v_1 \triangleleft e} \mathcal{F}_{v_1 \triangleleft e}^* \mathcal{F}_{v_1 \triangleleft e}$.*

Proof. Write down the block matrices representing δ and δ^* and compute. □

2 Operations on sheaves

2.1 Sheaf morphisms

When we define mathematical objects, we usually want to know how to define relationships between them. Often the objects themselves are less important than the maps we can build between them. Linear algebra is a prime example of the supremacy of maps over objects. Vector spaces, as objects, exist so that we can study linear transformations. Linear maps are far more interesting than vector spaces themselves. We can take this to an extreme and forget all the internal structure of our objects, looking instead at the structure of the relationships between different objects of the same type. This is the point of view of category theory, which serves as a deep organizing principle for much of mathematics. A category is formally a collection of objects together with morphisms between those objects.

So it is with sheaves. Sheaves have more structure than vector spaces, enough to be interesting on their own, but relationships between sheaves are still of paramount importance. This is why we define maps between sheaves. The first sorts of maps we will consider are between two sheaves over the same graph. All we need to deal with to define maps between sheaves on the same graph is linear algebra.

Definition 2.1 (Sheaf morphisms). Let \mathcal{F} and \mathcal{G} be two sheaves on a graph G . A *morphism* or *map* $\varphi : \mathcal{F} \rightarrow \mathcal{G}$ consists of a linear map $\varphi_v : \mathcal{F}(v) \rightarrow \mathcal{G}(v)$ for each vertex v of G and a linear map $\varphi_e : \mathcal{F}(e) \rightarrow \mathcal{G}(e)$ for each edge e , such that these linear maps commute with the restriction maps. That is, for each incident vertex-edge pair (v, e) , we have $\varphi_e \circ \mathcal{F}_{v \triangleleft e} = \mathcal{G}_{v \triangleleft e} \circ \varphi_v$.

This condition is easy to visualize in a diagram showing the relevant vector spaces and maps between them. The diagram *commutes* if any two paths along the arrows represent equal maps.

$$\begin{array}{ccc}
\mathcal{F}(v) & \xrightarrow{\varphi_v} & \mathcal{G}(v) \\
\downarrow \mathcal{F}_{v \triangleleft e} & & \downarrow \mathcal{G}_{v \triangleleft e} \\
\mathcal{F}(e) & \xrightarrow{\varphi_e} & \mathcal{G}(e)
\end{array}$$

A sheaf morphism can be represented by a large commuting diagram with many squares like the one above. The fact that the two sheaves are over the same graph means that the structure of a morphism is easy to define. The individual maps defining a sheaf morphism assemble into two maps $\varphi^0 : C^0(G; \mathcal{F}) \rightarrow C^0(G; \mathcal{G})$ and $\varphi^1 : C^1(G; \mathcal{F}) \rightarrow C^1(G; \mathcal{G})$, and the fact that the maps on stalks commute with restriction maps means that the amalgamated maps commute with the coboundary maps: $\delta^{\mathcal{G}} \circ \varphi^0 = \varphi^1 \circ \delta^{\mathcal{F}}$.

Because sheaf morphisms are made out of linear maps, they share a number of properties with them. Every linear map has a kernel, and so does every sheaf morphism. What may be surprising is that the kernel of a sheaf morphism is a sheaf. The stalks are easy to define: $(\ker \varphi)(v) = \ker \varphi_v$ and $(\ker \varphi)(e) = \ker \varphi_e$. We can restrict $\mathcal{F}_{v \triangleleft e}$ to $\ker \varphi_v$, getting a map $\mathcal{F}_{v \triangleleft e}|_{\ker \varphi_v} : \ker \varphi_e \rightarrow \mathcal{F}(e)$. If the image of this map lies in $\ker \varphi_e$, we have produced a restriction map $\mathcal{F}'_{v \triangleleft e} : (\ker \varphi)(v) \rightarrow (\ker \varphi)(e)$. Suppose $x \in \ker \varphi_v$. Then $\varphi_v x = 0$, so $\mathcal{G}_{v \triangleleft e} \varphi_v x = 0$. Because $\mathcal{G}_{v \triangleleft e} \circ \varphi_v = \varphi_e \circ \mathcal{F}_{v \triangleleft e}$, we know $\varphi_e \mathcal{F}_{v \triangleleft e} x = 0$. But this means that $\mathcal{F}_{v \triangleleft e} x \in \ker \varphi_e$, as required.

Remark. The sort of argument used above is known as a *diagram chase*. They can become quite complicated, and proofs of this kind are usually best communicated by drawing a commutative diagram, and pointing at various objects in the diagram while discussing the fate of an element as it is transferred through various maps.

A similar diagram chasing argument shows that sheaf morphisms can be composed just like linear maps. Composing the stalkwise maps of two sheaf morphisms yields another sheaf morphism.

For every sheaf \mathcal{F} , there is an *identity morphism* $\text{id} : \mathcal{F} \rightarrow \mathcal{F}$, which restricts to the identity map on each stalk. A sheaf morphism $\varphi : \mathcal{F} \rightarrow \mathcal{G}$ is an *isomorphism* if there exists a morphism $\psi : \mathcal{G} \rightarrow \mathcal{F}$ which is an inverse of φ ; that is, $\varphi \circ \psi = \text{id}$ and $\psi \circ \varphi = \text{id}$. One way to think of a sheaf isomorphism is as representing a “change of basis” of the sheaf. From an algebraic perspective, two isomorphic sheaves behave identically.

However, if we care about the inner product structure of the stalks of a sheaf, there is a finer distinction to be made. Invertible matrices preserve all algebraic properties of the vector spaces they map between, but they do not in general preserve the inner product structure. Similarly, sheaf isomorphisms do not automatically preserve the

inner product structure. A unitary sheaf isomorphism is one where the maps on stalks are unitary maps of inner product spaces; that is, they preserve the inner product.

Just as the space of linear maps from one vector space to another forms a vector space itself (consider the space of $m \times n$ matrices as the space of linear maps $\mathbb{R}^n \rightarrow \mathbb{R}^m$, for instance), so does the space of sheaf morphisms. The sum of two sheaf morphisms, taken by adding the component maps on stalks, is again a sheaf morphism. (Check this!) Scalar multiples of sheaf morphisms are again sheaf morphisms as well. We denote the space of sheaf morphisms $\mathcal{F} \rightarrow \mathcal{G}$ by $\text{Hom}(\mathcal{F}, \mathcal{G})$.

Example. Let \mathcal{F} be a sheaf on G , and consider the set of morphisms from the constant sheaf $\underline{\mathbb{R}}$ to \mathcal{F} . A linear map $\varphi_v : \mathbb{R} \rightarrow \mathcal{F}(v)$ is completely determined by $\varphi_v(1)$, so we may think of a morphism $\varphi : \underline{\mathbb{R}} \rightarrow \mathcal{F}$ as a collection of elements $x_v \in \mathcal{F}(v)$ and $x_e \in \mathcal{F}(e)$. The commutativity condition for being a sheaf morphism is that $\mathcal{F}_{v \triangleleft e} \circ \varphi_v = \varphi_e \circ \mathbb{R}_{v \triangleleft e}$. Since $\mathbb{R}_{v \triangleleft e}$ is the identity, we have $\mathcal{F}_{v \triangleleft e} x_v = x_e$. In other words, a morphism from $\underline{\mathbb{R}}$ to \mathcal{F} is the same as a choice of $x_v \in \mathcal{F}(v)$ and $x_e \in \mathcal{F}(e)$ so that $\mathcal{F}_{v \triangleleft e} x_v = x_e$ whenever $v \triangleleft e$. This is precisely the same information that we get from a global section of \mathcal{F} ! Therefore, we have

Proposition 2.1 (Global sections are morphisms). *If \mathcal{F} is a sheaf on G , there is a natural isomorphism between the vector spaces $\text{Hom}(\underline{\mathbb{R}}, \mathcal{F})$ and $\Gamma(G; \mathcal{F})$.*

Proposition 2.2 (Global sections are functorial). *Let $\varphi : \mathcal{F} \rightarrow \mathcal{G}$ be a sheaf morphism. This morphism induces a natural linear map $\Gamma(\varphi) : \Gamma(G; \mathcal{F}) \rightarrow \Gamma(G; \mathcal{G})$. Further, if $\psi : \mathcal{G} \rightarrow \mathcal{A}$ is another sheaf morphism, we have $\Gamma(\psi \circ \varphi) = \Gamma(\psi) \circ \Gamma(\varphi)$.*

Proof. Note that φ induces a map $\varphi^0 : C^0(G; \mathcal{F}) \rightarrow C^0(G; \mathcal{G})$ that commutes with the coboundaries of \mathcal{F} and \mathcal{G} . In particular, then, if $x \in \ker \delta^{\mathcal{F}}$, then $\delta^{\mathcal{G}} \varphi^0(x) = \varphi^1 \delta^{\mathcal{F}} x = 0$, so $\varphi^0(x) \in \ker \delta^{\mathcal{G}}$. This means that φ^0 induces a map $\Gamma(\varphi) : \Gamma(G; \mathcal{F}) \rightarrow \Gamma(G; \mathcal{G})$ by restriction to $\ker \delta^{\mathcal{F}}$. It is obvious that $\psi^0 \circ \varphi^0 = (\psi \circ \varphi)^0$, and since $\Gamma(\psi), \Gamma(\varphi)$, and $\Gamma(\psi \circ \varphi)$ are obtained by restriction of these maps, we get $\Gamma(\psi \circ \varphi) = \Gamma(\psi) \circ \Gamma(\varphi)$. \square

Remark. The operation of taking global sections takes a sheaf and produces a vector space. The fact that this operation behaves nicely with respect to sheaf morphisms and linear maps means that taking global sections is *functorial*. It translates information from the category of sheaves to the category of vector spaces in a way that respects relationships between sheaves.

2.2 Sheaves and graph homomorphisms

There are several different notions of graph homomorphism. The definition we will use is more permissive than some. For our purposes, a graph homomorphism f from $G = (V_G, E_G)$ to $H = (V_H, E_H)$ is given by a map $f_V : V_G \rightarrow V_H$ and a map $f_E : E_G \rightarrow E_H \cup V_H$ such that if v is incident to e , then either $f_V(v)$ is incident to $f_E(e)$ or $f_V(v) = f_E(e)$. This is an extension of the most common notion of graph homomorphism, which is a map on vertices that preserves the adjacency relation. They become equivalent if we stipulate that a vertex is adjacent to itself. However, the explicit treatment of edges in this notion of homomorphism is necessary for use with sheaves, since both edges and vertices carry data.

Remark. This notion of graph homomorphism is what we get when we consider graphs as *cell complexes* or *simplicial complexes*, and as such it extends to higher-dimensional structures than graphs.

If $f : G \rightarrow H$ is a graph homomorphism, we can take transfer sheaves on G or H across f to get sheaves on H or G . These are known as the pushforward and pullback operations.

Definition 2.2 (Pullback). Let $f : G \rightarrow H$ be a graph homomorphism, and let \mathcal{F} be a sheaf on H . The *pullback* of \mathcal{F} over f , $f^*\mathcal{F}$, is a graph on G with stalks $(f^*\mathcal{F})(\sigma) = \mathcal{F}(f(\sigma))$ and restriction maps $(f^*\mathcal{F})_{v \triangleleft e} = \mathcal{F}_{f(v) \triangleleft f(e)}$.

Proposition 2.3. *There exist morphisms $(f^*)^i : C^i(H; \mathcal{F}) \rightarrow C^i(G; f^*\mathcal{F})$ that commute with the coboundary maps of \mathcal{F} and $f^*\mathcal{F}$. That is, $\delta^{f^*\mathcal{F}} \circ (f^*)^0 = (f^*)^1 \circ \delta^{\mathcal{F}}$.*

Proof. For any vertex v of H , there is an obvious map $\mathcal{F}(v) \rightarrow f^*\mathcal{F}(v')$ whenever $f(v') = v$. We can assemble these into a map $\mathcal{F}(v) \rightarrow \bigoplus_{f(v')=v} f^*\mathcal{F}(v')$. Then combining these maps over vertices v of H , we get a map $(f^*)^0 : C^0(H; \mathcal{F}) = \bigoplus_{v \in H} \mathcal{F}(v) \rightarrow \bigoplus_{v \in H} \left(\bigoplus_{f(v')=v} f^*\mathcal{F}(v') \right) = \bigoplus_{v' \in G} f^*\mathcal{F}(v') = C^0(G; f^*\mathcal{F})$. A completely analogous argument furnishes a map $(f^*)^1 : C^1(H; \mathcal{F}) \rightarrow C^1(G; f^*\mathcal{F})$. Another, perhaps simpler, way to think of this map is as letting $((f^*)^i \chi)_\sigma = \chi_{f(\sigma)}$.

Commutativity will be satisfied if for every edge e of H , we have $(\delta^{f^*\mathcal{F}}(f^*)^0 \chi)_e = ((f^*)^1 \delta^{\mathcal{F}} \chi)_e$. The right hand side is equal to $(\delta^{\mathcal{F}} \chi)_{f(e)}$, while the left hand side is equal to

$$\begin{aligned} & (f^*\mathcal{F})_{h(e) \triangleleft e} ((f^*)^0 \chi)_{h(e)} - (f^*\mathcal{F})_{t(e) \triangleleft e} ((f^*)^0 \chi)_{t(e)} \\ &= \mathcal{F}_{f(h(e)) \triangleleft f(e)} \chi_{f(h(e))} - \mathcal{F}_{f(t(e)) \triangleleft f(e)} \chi_{f(t(e))} = (\delta^{\mathcal{F}} \chi)_{f(e)}. \end{aligned}$$

□

Definition 2.3 (Pushforward). Let $f : G \rightarrow H$ be a graph homomorphism, and let \mathcal{F} be a sheaf on G . The *pushforward* of \mathcal{F} over f , $f_*\mathcal{F}$, is a sheaf on H with stalks $f_*\mathcal{F}(\sigma) = \Gamma(f^{-1}(\sigma); \mathcal{F})$.

The restriction maps are slightly more complicated to specify. Given an element $x \in \Gamma(f^{-1}(v); \mathcal{F})$ and an edge e incident to v , every edge $e' \in f^{-1}(e)$ is incident to a unique vertex $v(e') \in f^{-1}(v)$. The local section x has a value $x_{v(e')}$ at each of these vertices, and hence we can define $(f_*\mathcal{F})_{v \triangleleft e}(x) = \sum_{e' \in f^{-1}(e)} \mathcal{F}_{v(e') \triangleleft e'}(x_{v(e')})$.

If a graph homomorphism sends edges only to edges, the pushforward is easier to define. In this case, $f_*\mathcal{F}(\sigma) = \bigoplus_{f(\tau)=\sigma} \mathcal{F}(\tau)$ and $(f_*\mathcal{F})_{v \triangleleft e} = \bigoplus_{f(\sigma)=v, f(\tau)=e} \mathcal{F}_{\sigma \triangleleft \tau}$.

There are two operations which combine two sheaves on the same graph into a single sheaf.

Definition 2.4 (Direct sum). If \mathcal{F} and \mathcal{G} are sheaves on a graph G , their *direct sum* $\mathcal{F} \oplus \mathcal{G}$ is the sheaf with stalks $(\mathcal{F} \oplus \mathcal{G})(v) = \mathcal{F}(v) \oplus \mathcal{G}(v)$, $(\mathcal{F} \oplus \mathcal{G})(e) = \mathcal{F}(e) \oplus \mathcal{G}(e)$ and restriction maps $(\mathcal{F} \oplus \mathcal{G})_{v \triangleleft e} = \mathcal{F}_{v \triangleleft e} \oplus \mathcal{G}_{v \triangleleft e}$.

Example. The zero sheaf $\underline{0}$, which assigns the zero vector space to each vertex and edge, is an identity for the direct sum. That is, for any sheaf \mathcal{F} , $\mathcal{F} \simeq \underline{0} \oplus \mathcal{F}$.

Definition 2.5 (Tensor product). If \mathcal{F} and \mathcal{G} are sheaves on a graph G , their *tensor product* $\mathcal{F} \otimes \mathcal{G}$ is the sheaf with stalks $(\mathcal{F} \otimes \mathcal{G})(v) = \mathcal{F}(v) \otimes \mathcal{G}(v)$, $(\mathcal{F} \otimes \mathcal{G})(e) = \mathcal{F}(e) \otimes \mathcal{G}(e)$ and restriction maps $(\mathcal{F} \otimes \mathcal{G})_{v \triangleleft e} = \mathcal{F}_{v \triangleleft e} \otimes \mathcal{G}_{v \triangleleft e}$.

One can naturally identify the coboundary map of $\mathcal{F} \oplus \mathcal{G}$ with the direct sum of the respective coboundary maps; the analogous statement does not hold for the coboundary map of $\mathcal{F} \otimes \mathcal{G}$. (This is because direct sums of vector spaces do not distribute over tensor products.)

Example. The constant sheaf $\underline{\mathbb{R}}$ is an identity of sorts for the tensor product. That is, for any sheaf \mathcal{F} , $\mathcal{F} \simeq \underline{\mathbb{R}} \otimes \mathcal{F}$.

2.3 Product graphs

There are several notions of the product of two graphs. Perhaps the most intuitive is the Cartesian product $G \square H$. This is the graph with vertex set $V(G) \times V(H)$, and an edge between (v_G, v_H) and (v'_G, v'_H) if either $v_G = v'_G$ and $v_H \sim v'_H$ or $v_G \sim v'_G$ and $v_H = v'_H$. This graph carries two projection homomorphisms $\pi_G : G \square H \rightarrow G$ and $\pi_H : G \square H \rightarrow H$.

If G and H carry sheaves \mathcal{F} and \mathcal{G} , there is a natural sheaf on $G \square H$ given by $\mathcal{F} \boxtimes \mathcal{G} = \pi_G^* \mathcal{F} \otimes \pi_H^* \mathcal{G}$. More concretely, $\mathcal{F} \boxtimes \mathcal{G}(v_G, v_H) = \mathcal{F}(v_G) \otimes \mathcal{G}(v_H)$, and the stalk associated with the edge between (v_G, v_H) and (v'_G, v_H) is $\mathcal{F}(v_G \sim v'_G) \otimes \mathcal{G}(v_H)$, and similarly for the edge between (v_G, v_H) and (v_G, v'_H) . The restriction map from (v_G, v_H) to $(v_G \sim v'_G, v_H)$ is $\mathcal{F}_{v_G \triangleleft (v_G \sim v'_G)} \otimes \text{Id}_{\mathcal{G}(v_H)}$.

If \mathcal{F} and \mathcal{G} are the constant sheaf $\underline{\mathbb{R}}$ on their respective graphs, then $\mathcal{F} \boxtimes \mathcal{G}$ is the constant sheaf $\underline{\mathbb{R}}$ on $G \square H$.

Remark. The notation $\mathcal{F} \boxtimes \mathcal{G}$ is adapted from the topological world, where the most important product is the Cartesian product, and this “outer product” of sheaves is frequently useful.

The tensor product $G \times H$ is the easiest to define: the vertex set is again $V(G) \times V(H)$, with an edge between (v_G, v_H) and (v'_G, v'_H) if $v_G \sim v'_G$ and $v_H \sim v'_H$. This graph also carries projection homomorphisms $\pi_G : G \times H \rightarrow G$ and $\pi_H : G \times H \rightarrow H$. As a result, we can define a sheaf $\mathcal{F} \times \mathcal{G}$ on $G \times H$ by $\mathcal{F} \times \mathcal{G} = \pi_G^* \mathcal{F} \otimes \pi_H^* \mathcal{G}$.

3 Spectral sheaf theory

The sheaf Laplacian as defined in section 1 is an invariant of sheaves on a labeled graph with specified stalks and orthonormal bases. Unlike the graph Laplacian, however, it is not a complete invariant. That is, there are non-isomorphic sheaves on the same graph which have the same sheaf Laplacian. A trivial reason for this is that the sheaf Laplacian does not record very much about the structure of the sheaf over the edges of the graph, but even more generally, there is too much room for redundant realizations of the same matrix structure.

Example. The two sheaves in Figure 3 both have sheaf Laplacian

$$\begin{bmatrix} 2 & -2 \\ -2 & 4 \end{bmatrix}$$

but are not isomorphic.

Remark. One important fact about sheaf Laplacians is that one may effectively delete a vertex or edge by assigning it the zero vector space. This is an extension of the way that one may delete edges from a graph by setting their weights to zero.

Although the sheaf Laplacian is not a complete invariant of sheaves, it does convey a lot of information about the sheaf. One obvious example is whether the sheaf has




Figure 3: Two nonisomorphic sheaves with identical sheaf Laplacians

any global sections, since $\ker L_{\mathcal{F}} = \ker \delta$. The smallest nontrivial eigenvalue of a sheaf Laplacian gives us information about how close the sheaf is to having more global sections.

The spectrum of the sheaf Laplacian interacts in interesting ways with the sheaf operations described above. These often generalize results known for the spectra of graph Laplacians.

3.1 Basic facts about sheaf Laplacians

Proposition 3.1. *Let \underline{V} be the constant sheaf on a graph G . The sheaf Laplacian of \underline{V} , with respect to an orthonormal basis of V , is $L_G \otimes I$, where L_G is the graph Laplacian of G .*

A common question in spectral graph theory is how the spectrum of a graph changes as we manipulate the graph. These tend to be the easiest results to extend to spectral sheaf theory, since the analogous theorem statements are usually immediately evident.

A useful concept when studying the spectra of related matrices is an eigenvalue interlacing.

Definition 3.1. Let A, B be $n \times n$ matrices with real spectra, and let $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ be the eigenvalues of A and $\mu_1 \leq \mu_2 \leq \dots \leq \mu_n$ be the eigenvalues of B . We say the eigenvalues of A are **(p, q)-interlaced** with the eigenvalues of B if for all k , $\lambda_{k-p} \leq \mu_k \leq \lambda_{k+q}$. (We let $\lambda_k = \lambda_1$ for $k < 1$ and $\lambda_k = \lambda_n$ for $k > n$.)

We now show that the eigenvalues of sheaf Laplacians are interlaced after deleting a subgraph.

Theorem 3.1 (Eigenvalue interlacing for the sheaf Laplacian). *Let \mathcal{F} be a sheaf on a graph G , C a collection of edges of G , and let $H = G \setminus C$ denote G with the edges in C removed.*

Let \mathcal{G} be the sheaf with the same vertex stalks as \mathcal{F} but with all edge stalks over edges not in C set to zero. Then the eigenvalues of L_H are $(t, 0)$ -interlaced with the eigenvalues of L_G , where $t = \text{codim } H^0(G; \mathcal{G}) = \dim C^0(G; \mathcal{F}) - \dim H^0(G; \mathcal{G})$.

Proof. This is a standard sort of argument in spectral graph theory. We use Rayleigh quotients to bound eigenvalues in terms of the quadratic form defined by the Laplacian. Notice that $L_H = L_G - L_C$, where L_C is the Laplacian of \mathcal{G} . By the Courant-Fischer theorem, we have

$$\begin{aligned} \mu_k &= \min_{\dim Y=k} \left(\max_{\mathbf{y} \in Y, \mathbf{y} \neq 0} \frac{\langle \mathbf{y}, L_G \mathbf{y} \rangle - \langle \mathbf{y}, L_C \mathbf{y} \rangle}{\langle \mathbf{y}, \mathbf{y} \rangle} \right) \\ &\geq \min_{\dim Y=k} \left(\max_{\mathbf{y} \in Y \cap H^0(G; \mathcal{G}), \mathbf{y} \neq 0} \frac{\langle \mathbf{y}, L_G \mathbf{y} \rangle}{\langle \mathbf{y}, \mathbf{y} \rangle} \right) \\ &\geq \min_{\dim Y=k-t} \left(\max_{\mathbf{y} \in Y, \mathbf{y} \neq 0} \frac{\langle \mathbf{y}, L_G \mathbf{y} \rangle}{\langle \mathbf{y}, \mathbf{y} \rangle} \right) = \lambda_{k-t} \end{aligned}$$

and

$$\begin{aligned} \lambda_k &= \min_{\dim Y=k} \left(\max_{\mathbf{y} \in Y, \mathbf{y} \neq 0} \frac{\langle \mathbf{y}, L_G \mathbf{y} \rangle}{\langle \mathbf{y}, \mathbf{y} \rangle} \right) \\ &\geq \min_{\dim Y=k} \left(\max_{\mathbf{y} \in Y, \mathbf{y} \neq 0} \frac{\langle \mathbf{y}, L_G \mathbf{y} \rangle - \langle \mathbf{y}, L_C \mathbf{y} \rangle}{\langle \mathbf{y}, \mathbf{y} \rangle} \right) = \mu_k. \end{aligned}$$

□

Unitary sheaf isomorphisms preserve the spectrum of the Laplacian. To see this, suppose that $\varphi : \mathcal{F} \rightarrow \mathcal{G}$ is a sheaf morphism such that φ^0 and φ^1 are unitary maps. Then because $\varphi^1 \delta^{\mathcal{F}} = \delta^{\mathcal{G}} \varphi^0$, we have $L_{\mathcal{F}} = (\delta^{\mathcal{F}})^*(\varphi^1)^* \varphi^1 \delta^{\mathcal{F}} = (\varphi^0)^*(\delta^{\mathcal{G}})^* \delta^{\mathcal{G}} \varphi^0 = (\varphi^0)^* L_{\mathcal{G}} \varphi^0$. Since φ^0 is unitary, the two matrices have the same spectrum. If the components of φ are just unitary maps, the spectrum of $L_{\mathcal{F}}$ is included in the spectrum of $L_{\mathcal{G}}$.

Definition 3.2 (Covering map). A graph morphism $f : G \rightarrow H$ is called a *covering map* if for every $v \in H$, the number of vertices w with $f(w) = v$ is constant, there are no edges that map to vertices, and the number of edges incident to $w \in G$ is the same as the number of edges incident to $f(w) \in H$.

TODO: Picture of a covering map

Proposition 3.2. *Let $f : G \rightarrow H$ be a covering map of graphs. If \mathcal{F} is a sheaf on H , the spectrum of $L_{f^* \mathcal{F}}$ contains the spectrum of $L_{\mathcal{F}}$.*

Proof. Consider the map $f^* : C^i(H; \mathcal{F}) \rightarrow C^i(G; f^*\mathcal{F})$ given by $(f^*\chi)_v = \chi_{f(v)}$. If χ is an eigenvector of $L_{\mathcal{F}}$ with eigenvalue λ , we have

$$\begin{aligned} (L_{f^*\mathcal{F}}f^*\chi)_v &= \sum_{v \triangleleft e, w \triangleleft e} ((f^*\mathcal{F})_{ve}^*(f^*\mathcal{F})_{ve}(f^*\chi)_v - (f^*\mathcal{F})_{ve}^*(f^*\mathcal{F})_{we}f^*\chi_w) \\ &= \sum_{f(v) \triangleleft e, f(w) \triangleleft e} \mathcal{F}_{f(v)e}^* \mathcal{F}_{f(v)e} \chi_{f(v)} - \mathcal{F}_{f(v)e}^* \mathcal{F}_{f(w)e} \chi_{f(w)} = (L_{\mathcal{F}}\chi)_{f(v)} = \lambda \chi_{f(v)} = \lambda (f^*\chi)_v. \end{aligned}$$

Since every eigenvector of $L_{\mathcal{F}}$ produces a corresponding eigenvector of $L_{f^*\mathcal{F}}$ with the same eigenvalue, the spectrum of $L_{f^*\mathcal{F}}$ contains the spectrum of $L_{\mathcal{F}}$. \square

A number of other interesting results related to sheaves and their Laplacians may be found in [3]. Among these are results on:

- spectra of pushforward sheaves
- sheaves on Cartesian product graphs
- effective resistance on sheaves
- spectral sparsification of sheaves
- harmonic functions on sheaves.

4 Directions for applications

4.1 Diffusion

The existence of a Laplacian immediately suggests consideration of its associated differential equation. This is the Laplacian flow, or diffusion

$$\dot{\chi} = -L_{\mathcal{F}}\chi.$$

Since $L_{\mathcal{F}}$ is positive semidefinite with kernel equal to $H^0(G; \mathcal{F})$, the trajectories of this dynamical system converge to global sections of \mathcal{F} .

Graph Laplacian-based diffusions are often used as a building block for the construction or study of network dynamics. Sheaf Laplacians can add richness to these models while remaining analytically tractable. Opinion dynamics in social networks offer an illustrative example. Perhaps the simplest model of opinion dynamics posits a social network described by a graph G , where each agent has a one-dimensional space of opinions. The dynamics are described by the Laplacian flow on G , and eventually converge to consensus.

Replacing G with a sheaf on G allows us to model various more interesting aspects of the system:

- We can add extra dimensions to the opinion space by increasing the dimension of the vertex and edge stalks
- We can implement links that force opinions apart for certain agents by changing the sign of one of the restriction maps on an edge
- We can model the difference between one's privately-held opinion and the one communicated to others by changing edge stalks and restriction maps.

While these additions are still limited by the fact that the overall dynamics are linear, they introduce interesting behavior in such a system, and can all be analyzed within the sheaf framework. Further, the linear dynamics can be used as a building block for more sophisticated models.

4.2 Distributed consensus

Diffusion on graphs is also used to study consensus dynamics on constructed networks. Suppose we have a collection of agents connected according to the edges in a graph, and each has some internal state in some vector space V . We want them to communicate over the edges of the graph in order to come to an agreement on a value. We can represent this condition as saying we wish the agents to reach a global section of the constant sheaf \underline{V} . This can be achieved by having them follow the Laplacian flow associated to the sheaf.

This requires each agent to communicate its full state to all of its neighbors. This is perhaps an onerous requirement, but is not necessary. Sheaf theory will help us reduce the bandwidth requirements.

Definition 4.1 (Approximation to a sheaf). Let G be a graph, and let \mathcal{G} be a sheaf on G . We say that a sheaf \mathcal{F} on G is an *approximation to* \mathcal{G} if there exists a morphism $\alpha : \mathcal{F} \rightarrow \mathcal{G}$ which is an isomorphism on vertex stalks, and which induces an isomorphism $\Gamma(G; \mathcal{G}) \rightarrow \Gamma(G; \mathcal{F})$.

If V is a vector space, we denote the constant sheaf with stalk V by \underline{V} , and say that \mathcal{F} is an *approximation to the constant sheaf* if \mathcal{F} is an approximation to \underline{V} .

Proposition 4.1. *If \mathcal{F} is an approximation to \underline{V} , then it is isomorphic to a sheaf with vertex stalks V where the restriction maps $\mathcal{F}_{ve} : V \rightarrow \mathcal{F}(e)$ and $\mathcal{F}_{v'e} : V \rightarrow \mathcal{F}(e)$ are equal.*

Proof. Note that because $\alpha : \underline{V} \rightarrow \mathcal{F}$ is an isomorphism on vertex stalks, \mathcal{F} is clearly isomorphic to a sheaf with vertex stalks V . For every edge e we have the diagram

$$\begin{array}{ccc} V & \xrightarrow{1} & V \\ \downarrow 1 & & \downarrow \mathcal{F}_{ve} \\ V & \xrightarrow{\alpha_e} & \mathcal{F}(e) \\ 1 \uparrow & & \uparrow \mathcal{F}_{we} \\ V & \xrightarrow{1} & V \end{array},$$

and the only way it can commute is if $\mathcal{F}_{ve} = \mathcal{F}_{we} = \alpha_e$.

□

The proof of this proposition shows that specifying an approximation to \underline{V} is the same as specifying a morphism $\alpha_e : V \rightarrow W_e$ for each edge e of G . Further, in order to produce an approximation to \underline{V} , the α_e must assemble to a map $\alpha : C^1(G; \underline{V}) \rightarrow C^1(G; \mathcal{F}) = \bigoplus_{e \in E} W_e$ such that $\ker(\alpha \circ \delta_{\underline{V}}) = \ker \delta_{\underline{V}}$. This holds if $\ker \alpha$ is contained in a complement to $\text{im } \delta$; equivalently, the projection map $\pi : C^1(G; \underline{V}) \rightarrow H^1(G; \underline{V})$ must be an isomorphism when restricted to $\ker \alpha$.

This inspires a way to construct an approximation to the constant sheaf. Choose a subspace K_e of V for each edge e of G and define α_e to be the projection map $V \rightarrow V/K_e$. If $\bigoplus_{e \in E} K_e$ has the same dimension in $H^1(G; \underline{V})$ as in $C^1(G; \underline{V})$, then $\alpha = \bigoplus_{e \in E} \alpha_e$ defines the edge maps giving an approximation to \underline{V} . (The vertex maps may be taken to be the identity.)

A spectrally good approximation to the constant sheaf (satisfying a few other conditions) is analogous to an expander graph. It is possible that expander sheaves would allow for implementations of faster consensus algorithms for high-dimensional data on graphs by reducing the amount of communication needed to ensure convergence.

4.3 Learning Sheaves

From a network science perspective, it may be interesting to try to find a sheaf for which a given collection of signals is close to being sections. That is, we are given a number of vectors $x_i \in C^0(G; \mathcal{F})$, and wish to find both a graph G and a sheaf \mathcal{F} on G so that $\sum_i x_i^T L_{\mathcal{F}} x_i$ is small. It turns out that the space of sheaf Laplacians is a convex cone, so we can cast this as a convex optimization problem. Seeking to minimize $\sum_i x_i^T L_{\mathcal{F}} x_i$

alone leads to a minimum at $L_{\mathcal{F}} = 0$, so we need to add regularization terms. Details can be found in [4].

What does the learned $L_{\mathcal{F}}$ tell us? We can extract an underlying graph structure from the sparsity pattern of the Laplacian. These are connections implied by the data. Each edge further includes a description of the tendencies of the relationships in the data: an edge “pushes” the data on its incident vertices to lie in some algebraic relationship.

5 Where to go from here

5.1 Beyond graphs

In algebraic topology, graphs have higher-dimensional generalizations known as simplicial complexes and cell complexes. Graphs are one-dimensional cell complexes, and simple graphs are one-dimensional simplicial complexes. Sheaves can be defined on cell complexes and simplicial complexes by adding stalks for higher-dimensional cells and restriction maps for each incidence relation between cells. These give us a way to encode higher-order consistency conditions, but these do not appear in quite the way one might expect. The coboundary map extends to a sequence of coboundary maps between cochain spaces of adjacent cell degrees, but the significance of the kernel of higher coboundary maps is less obvious.

Even more generally, it is possible to define a sheaf on a partially-ordered set; this is the approach taken by Michael Robinson in his work on sheaves and data analytics. This case subsumes the cell complex case, since the face relations of a cell complex form a poset.

There are versions of sheaves for general topological spaces and even more exotic structures like sites, but thus far these have been largely irrelevant to engineering applications.

5.2 Beyond vector spaces

There is nothing in the definition of a sheaf that requires the stalks actually be vector spaces. One may construct a sheaf *valued in* any number of other categories. For instance, the stalks might be simply sets, and the restriction maps just functions. Or they might be groups, and the restriction maps group homomorphisms. The difficulty here is that the tools we have developed for sheaves of vector spaces do not have immediate

analogues when the stalks are not vector spaces. One situation that allows some of the machinery developed for vector spaces is the case of sheaves of semimodules over a semiring. (A semiring is the analogue of a ring without the requirement that every element have an additive inverse.) These give a natural way to model nonnegativity constraints and directionality.

5.3 Turning things around

Some relationships we would like to model are more naturally expressed with maps going the opposite way. Rather than restriction maps going from vertex stalks to edge stalks, we want to consider extension maps going from edge stalks to vertex stalks. In category theory, it is common to take the *dual* of a construction by reversing the direction of all relevant morphisms. When we do this to a sheaf, we get what is called a *cosheaf*.

Cosheaves are somewhat less intuitive than sheaves. Rather than global sections, they have *cosections*, which are assignments to vertex stalks modulo an equivalence relation given by the edges. Similarly, rather than *cohomology*, they have homology, and the cosections of a cosheaf are precisely the elements of the degree zero homology space H_0 .

5.4 Further reading

Most other resources about sheaves have a heavier topological and algebraic emphasis, and more prerequisites. The goal of these notes is to prepare the reader interested in applications to tackle the more sophisticated treatments.

Elementary Applied Topology, by Robert Ghrist [2]. This book serves as an introduction to algebraic topology and its applications to a wide variety of problems. The penultimate chapter discusses sheaves, but is not accessible without a good deal of the previous material in the book.

Topological Signal Processing, by Michael Robinson [5]. This book is written from a more explicit engineering point of view, with examples of specific applications. Not everything in here is relevant to the study of sheaves on graphs, though, and it focuses less on networks. Robinson also has slides and video from a seminar on sheaves and data he presented for DARPA: **Tutorial on Sheaves in Data Analytics**.

Sheaves, Cosheaves, and Applications, by Justin Curry [1]. Justin's thesis is a perennial source of new insights, but is written from the most abstract perspective of

any of these resources. It is not likely to be a fruitful read for anyone not comfortable with the material in Elementary Applied Topology.

Toward a Spectral Theory of Cellular Sheaves, by Jakob Hansen and Robert Ghrist [3]. This is an expository paper intended for an audience in applied topology. It contains more general versions of most of the results discussed in this introduction.

References

- [1] J. Curry. *Sheaves, Cosheaves, and Applications*. PhD thesis, University of Pennsylvania, 2014.
- [2] R. Ghrist. *Elementary Applied Topology*. CreateSpace, <https://www.math.upenn.edu/~ghrist/notes.html>, 2014.
- [3] J. Hansen and R. Ghrist. Toward a Spectral Theory of Cellular Sheaves. *arXiv:1808.01513 [math]*, Aug. 2018.
- [4] J. Hansen and R. Ghrist. Learning Sheaf Laplacians from Smooth Signals. In *Proceedings of ICASSP*, 2019.
- [5] M. Robinson. *Topological Signal Processing*. Springer-Verlag Berlin Heidelberg, 2014.